

Remarks

The above Amendments and these Remarks are in reply to the Office action mailed June 10, 2005.

Currently, claims 1-44 are pending. Applicants have amended claims 1, 3-14, 16, 21-25, 27, 33-34, 38, and 42 and cancelled claims 2 and 26. Applicants respectfully request reconsideration of claims 1, 3-25, and 27-44.

Applicants have amended the specification at page 1 to include the application serial numbers for the cross-referenced applications.

I. Objection to Claims 3, 15, 17, 37, and 44

Claims 3, 15, 17, 37 and 44 were objected to because “it is not clear under what time frame is ‘the number of packets’ collected.”

Claim 3 has been amended to recite “wherein the step of providing quality of service includes controlling the number of packets from the at least one initiator to the at least one storage device during a period of time.”

Claims 15, 17, and 44 recite, “determining if a buffer includes a number of packets beyond a specified threshold.” Claim 37 recites, “when the buffer includes a number of packets from the initiator that exceeds a specified threshold.” It is respectfully submitted that these phrases are clear, unambiguous, and in accordance with 35 U.S.C. § 112, ¶¶ 1 and 2. As understood by those of ordinary skill in the art, a buffer is a temporary storage mechanism. In these claims, the buffer temporarily stores “packets.” The claims recite determining if the buffer “includes a number of packets beyond a specified threshold,” or “when the buffer includes a number of packets from the initiator that exceeds a specified threshold.” It is not necessary to state any time frame as it is clear that the recited limitations are not based on any time frame; rather, they are based on the actual “number of packets” in the buffer at any given time. Because it is clear from the unambiguous claim language that the limitations are not based on the number of packets in the buffer collected over some time period, Applicants assert that the objection is improper and respectfully request its withdrawal.

II. Objection to Claims 9-15, and 38-44

Claims 9-15 and 38-44 were objected to because “the term ‘the actual bandwidth’ appears to lack antecedent basis in claims 9, 14-15, 38 and 44.” Independent claims 9 and 38 have been amended to recite “an actual bandwidth.” Accordingly, it is respectfully submitted that the antecedent basis problem of claims 9, 14-15, 38, and 44 has been corrected. Withdrawal of the objection to claims 9-15 and 38-44 is respectfully requested.

III. Rejection of Claims 1-2, 25-27, 30, 33, and 35-37

Claims 1-2, 25-27, 30, 33 and 35-37 were rejected under 35 U.S.C. § 102(e) as being anticipated by U.S. 20020194324 (*Guha*). Claims 2 and 26 have been cancelled. Because *Guha* fails to disclose each limitation of claims 1, 25, 27, 30, 33, and 35-37, Applicants assert that these claims are patentable over the cited art.

An important feature of embodiments of Applicants’ invention is a “storage switch” that can be in direct communication with initiators (e.g., servers) and targets (e.g., storage devices or subsystems) within a storage area network so that it can “enforce Quality of Service” (QoS) for connections between initiators and targets.” *Applicants’ Specification*, ¶ 0015; *see also Figure 2*, ¶ 0018-0019. Applicants’ “storage switch,” a single device, can provision a virtual target (¶ 0075-0078), provision an initiator connection (¶ 0079-0082), and provision a user domain (¶ 0083-0087). A Quality of Service policy maintained in the switch can then be generally defined by the “virtual target (as provisioned), the initiator connection (as provisioned), and the User Domain.” *Id. at* ¶ 0089.

Applicants’ switch-based Quality of Service enforcement for a storage area network is novel. For example, “in a conventional communications network (as opposed to a storage network), QoS is used to ensure that users get the percentage of data bandwidth for a connection that they paid for.” *Id. at* ¶ 0108. This QoS in a conventional network “is enforced by prioritizing the switching traffic even at the expense of dropping packets.” *Id.* Applicants explain the uniqueness of providing QoS within a storage area network where “dropping a request in a storage system is unacceptable, unlike conventional network communication system, where a request may include one or more packets.” *Id. at* ¶ 0108.

Claims 1 and 2

Amended claim 1 captures such a concept, reciting:

the storage network including at least one initiator, at least one storage device, and a storage switch in communication with the at least one initiator and the at least one storage device, the method comprising:

providing, by the storage switch, quality of service to the at least one initiator for accessing the at least one storage device in the storage network.

Emphasis added.

The prior art of record fails to disclose such a concept as claimed. Specifically, the prior art fails to disclose “a storage switch” that is in communication with an “initiator” and a “storage device” so that quality of service can be provided “by the storage switch” to “the at least one initiator for accessing the at least one storage device,” as recited in claim 1 (*emphasis added*).

Guha’s system is comprised of multiple components. Looking at Figures 4 and 6, for example, it is seen that *Guha* includes a QoS Enforcer 34, a layer 4 switch, servers 39, 41, 40, content controller 36, a SAN switch 42 and storage devices 44. Neither of the switches 34 or 42 provide “quality of service,” as recited in claim 1.

Guha provides a specific component – a “QoS enforcer” (denoted as numeral 34 in FIGs. 3-6) – within his distributed content management system to provide quality of service. Thus, *Guha* fails to disclose “providing, by the storage switch, quality of service,” as recited in claim 1 (*emphasis added*). Quite simply, *Guha* teaches a separate component to provide quality of service, not a storage switch.

To the extent the Examiner regards the “layer 4 switch” 38 as providing quality of service or forming the “QoS enforcer” 34, *Guha* still fails to disclose providing quality of service by a “storage switch,” or by any switch that is “in communication with the at least one initiator and the at least one storage device.” *Emphasis added.* The “layer 4 switch” 38 of *Guha* is not a storage switch. As understood by those of ordinary skill in the art, a “layer 4 switch” is a standard network routing device that operates at the transport layer of TCP/IP, for example. *See, e.g., Layer 4 Switching: Unraveling the Vendorspeak*, May 3, 1999 (<http://www.networkcomputing.com/1009/1009ws1.html>). They do not operate at the storage level, nor do they provide storage functionality as a storage switch. This is clear from *Guha*’s own disclosure, which clearly illustrates in Figures 4 and 6 that the “layer 4 switch” is not part of the storage area network.

Moreover, the “layer 4 switch” 38 is not “in communication with the at least one initiator and the at least one storage device,” as claim 1 recites. As Figures 4 and 6 illustrate, the “layer 4 switch” is in communication with servers 39, 41, 40, not the storage devices 44. The servers 39-41 are depicted differently than storage devices 44, clearly illustrating that the servers are not operating as storage devices in *Guha*’s system. The storage devices 44 are connected to a different switch, SAN switch 42. Thus, they are not in communication with the layer 4 switch.

This distinction is important. As discussed above, the concept of Applicants’ invention as claimed is providing “by the storage switch, quality of service” in a storage network. Providing quality of service in a storage network is not the same as providing quality of service in a conventional network system. Applicants explain that, “dropping a request in a storage system is unacceptable, unlike conventional network communication system, where a request may include one or more packets.” *Id.* at ¶ 0108. Applicants claim providing quality of service “by the storage switch” which is “in communication with the at least one initiator and the at least one storage device” to meet the unique requirements of storage networks. *Guha* appears directed to quality of service at this “conventional communications network” level. Specifically, *Guha* teaches to drop requests, stating “if the QoS that can be provided is not high and the remote load of the architecture needed to comply with the request is high, then the request is delayed or dropped.” *Id.* at ¶ 0073 (*emphasis added*). Accordingly, *Guha* does not use a “storage switch” to provide quality of service nor does he provide the quality of service by a switch that “is in communication with the at least one initiator and the at least one storage device,” as recited in claim 1.

It should also be noted that switch 42 does not provide QoS in *Guha*. *Guha*’s quality of service is provided by the QoS Enforcer, shown at the front end of the systems shown in Figures 4 and 6. Switch 42 does not provide this quality of service.

The Examiner further references paragraph 15 of *Guha* as disclosing providing quality of service by a switch. Applicants respectfully disagree as paragraph 15 supports no such proposition. Paragraph 15 recites that:

In the system of the present invention each data storage center may further include: at least one server device which communicates with the QoS enforcer; a network switch which communicates with the at least one server device; and at least one storage device which communicates with the SAN switch.

This passage clearly shows that quality of service is provided by the QoS enforcer, not any of the switches. Moreover, this passage clearly shows that the network switch, which communicates with the QoS enforcer is in communication with the at least one server device, but not with the storage device. The storage device is in communication with a different switch, the SAN switch, which does not provide quality of service functionality.

Because *Guha* fails to disclose each of the limitations of claim 1, as highlighted above, Applicants assert that claim 1 is patentable over the cited art.

Claims 25-27

Claim 25 recites:

A switch for use in a storage network, the switch comprising:
a port to be coupled to an external device, wherein the external device includes at least one of an initiator and a storage device;
a bandwidth controller, the bandwidth controller including a processor, a traffic manager, and a buffer.

Guha does not disclose a switch for use in a storage network that includes a “bandwidth controller including a processor, traffic manager, and a buffer,” as recited in claim 25.

It is important to first note that the elements recited in claim 25 are all part of a single “switch.” Applicants point this out as the *Office Action* relies on multiple discrete devices for the assertion that *Guha* discloses each element of the claimed “switch.” Applicants have amended claim 25 to make it clear that the recited elements are all part of the “switch” and not merely part of the “storage network” in which the switch is for use.

As recited in claim 25, the switch includes a “bandwidth controller.” The “bandwidth controller” includes both “a processor” and a “traffic manager.” The “processor” and “traffic manager” are individually recited elements that are both part of the “switch.” *Guha* does not disclose a switch containing these elements.

The Examiner asserts that elements 34-40 of Figure 4 in *Gupta* disclose a “bandwidth controller.” The Examiner further asserts that elements 39-40 form the “processor” of the “bandwidth controller” and that element 34 forms the “traffic manager” of the “bandwidth controller.” Each element of *Guha* cited by the Examiner is a discrete individual device. Hence,

they are not all part of a “switch” as claim 25 recites. In fact, none of them are part of the switch included within this set of cited elements.

Guha discloses a “layer 4 switch” 38 which is coupled to a “QoS Enforcer 34” on one end and to “servers 39-41” on the other end. *See e.g., Figure 4.* The Examiner asserts that these “servers” form the “processor” of the “bandwidth controller” in the “switch” recited in claim 25. However, these servers are separate and distinct from the “layer 4 switch 38.” This is plainly clear as shown in Figure 4. Because they are individual devices apart from “switch 38,” they are not a “processor” of a “bandwidth controller” contained in a “switch,” as recited in claim 25.

The Examiner asserts that the “QoS Enforcer 34” is the “traffic manager” of the “bandwidth controller” in the “switch” recited in claim 25. Like the servers, “QoS Enforcer 34” is shown as a discrete component in *Guha*. Because the “QoS Enforcer” is an individual component apart from “layer 4 switch 38,” it is not a “traffic manager” of a “bandwidth controller” contained in a “switch,” as recited in claim 25.

In sum, *Guha* does not disclose a “switch” that includes the elements recited in claim 25. Rather, *Guha* shows a system including many individual components that are not part of a “switch.” Accordingly, *Guha* fails to disclose a “switch” as recited in claim 25. Claims 26-27 each ultimately depend from claim 25 and therefore, should be patentable for at least the same reasons.

Claim 30

Claim 30 recites:

A switch, including:
a storage processor, including a request controller;
a traffic manager in communication with the storage processor;
a buffer in communication with the traffic manager;
wherein if a specified threshold in the buffer is reached, the traffic manager is designed to activate the request controller.

For the same reasons set forth with respect to claim 25, Applicants assert that *Guha* fails to disclose each limitation of claim 30. For example, *Guha* fails to disclose a “switch” that includes “a storage processor” and “a traffic manager.”

Applicants further assert that *Guha* fails to disclose the additional limitation of “wherein if a specified threshold in the buffer is reached, the traffic manager is designed to activate the request

controller,” recited in claim 30. The Examiner cites to paragraphs 0068-0069 of *Guha* asserting that a “buffer must have existed in order to performing the tasks described in paragraphs 0068-0069.” However, a close look at these paragraphs reveals that they do no support the disclosure of this limitation of claim 30. For example, *Guha* explains that “based on traffic levels observed and communicated by the QoS Enforcer 34 to the content controller 36, additional resources at the server and storage levels can be reassigned in the content pool to improve I/O access and the SLA needs of the content requests.” *Guha*, ¶ 0069 (*emphasis added*). There is nothing in this disclosure to support a teaching or even a suggestion of “if a specified threshold in the buffer is reached, the traffic manager is designed to activate the request controller,” as recited in claim 30. The cited passage merely states that traffic levels are “observed and communicated.” *Guha* does not explain how they are observed. Moreover, *Guha* simply explains that resources can be reassigned in response to the observed levels. There is no disclosure of a “traffic manager” that “activate[s] the request controller” based on a specified threshold in the buffer being reached. Accordingly, *Guha* cannot fairly be said to teach or suggest a “traffic manager” activating a “request controller” based on “whether a specified threshold in the buffer is reached,” as recited in claim 30.

Because *Guha* fails to disclose each limitation of claim 30, Applicants assert that claim 30 is patentable over the cited art.

Claims 33 and 35

Claim 33 recites:

A storage switch for use in a storage network comprising:
a first port to be coupled to at least one initiator;
a second port to be coupled to at least one storage device; and
means for providing quality of service for a connection from the at least one initiator to the at least one storage device in the storage network.

As set forth with respect to claim 1, *Guha* fails to disclose a “storage switch” that is to be coupled to “at least one initiator” and “at least one storage device” and that provides “quality of service for a connection from the at least one initiator to the at least one storage device in the storage network.” Accordingly, for at least the same reasons as set forth above, Applicants assert that claim 33 is patentable over the cited art. Claim 35 depends from claim 33 and therefore, should be patentable for at least the same reasons.

Claim 37

Claim 37 recites a storage network that includes “a switch” including “a traffic manager in communication with a buffer; wherein when the buffer includes a number of packets from the initiator that exceeds a specified threshold, then the switch is designed to notify the initiator to reduce a number of concurrent requests.” As shown with respect to claim 33 which recites similar limitations, *Guha* does not disclose a “switch” having each of these elements. Accordingly, for at least the same reasons as set forth above, Applicants assert that claim 37 is patentable over the cited art.

IV. Rejection of Claims 3-24, 29, 31-32, 34, and 38-44 under 35 U.S.C. § 103(a)

Claims 3-24, 29, 31-32, 34 and 38-44 were rejected under 35 U.S.C. 103(a) as being unpatentable over *Guha*, as applied to claims 1-2, 25-27, 30, 33 and 35-37, further in view of U.S. Pat. No. 5,719,854 (*Choudhury*). Because *Guha* and *Choudhury*, either alone or in combination, fail to teach or suggest each of the limitations of claims 3-24, 29, 31-32, 34 and 38-44, Applicants assert that these claims are patentable over the cited art.

Claims 9-15

Among other limitations, claim 9 recites:

A method for use in a storage network, the storage network including at least one initiator, at least one storage device, and at least one storage switch, wherein the at least one initiator and the at least one storage device are both in communication with the storage switch, the method comprising:

guaranteeing, by the storage switch, a minimum bandwidth to the at least one initiator to access the at least one storage device in the storage network;

measuring, by the storage switch, an actual bandwidth utilized by the initiator, where the actual bandwidth is measured by a number of requests per second times an average size of requests from the at least one initiator. *Emphasis added.*

As detailed above with respect to claim 1, *Guha* fails to disclose a “storage switch” that is in communication with an initiator and storage device and that provides quality of service to the initiator. *Guha* further fails to teach or suggest such a “storage switch” that guarantees “a minimum bandwidth to the at least one initiator to access the at least one storage device in the storage

network,” and measures “the actual bandwidth utilized by the at least one initiator,” as claim 9 recites.

In *Guha*, a separate component – the “QoS enforcer” – provides the quality of service functionality and any bandwidth control. It is not provided, and there is nothing to suggest that it is provided, by a “storage switch,” as claim 9 recites. The “QoS enforcer” “monitors content requests that arrive at the data center and controls the entrance of all traffic.” *Guha*, ¶ 0039 (*emphasis added*). In order to control the entrance of traffic to the data center itself, *Guha* positions the “QoS enforcer” at the head of each individual data center. Were the QoS enforcer positioned other than at the head of the data center, it would appear that the purpose of controlling “the entrance of all traffic” would be frustrated. Thus, *Guha* contains nothing to suggest the claimed operations at the storage level by a “storage switch,” as recited in claim 9.

To the extent the Examiner regards the “layer 4 switch” 38 as guaranteeing “a minimum bandwidth” or forming the “QoS enforcer” 34, *Guha* still fails to teach or suggest such a function by a “storage switch,” or by any switch that is “in communication with the at least one initiator and the at least one storage device,” as claim 9 recites (*emphasis added*). The “layer 4 switch” of *Guha* is not a “storage switch,” nor is it in communication with “at least one storage device” as claim 9 recites. The “layer 4 switch” is a network routing device that is communication with servers 39, 40, and 41. It is not in communication with storage devices 44.

Lastly, *Guha* contains no teaching or suggestion that the SAN switch 42 disclosed therein performs any quality of service functionality such as guaranteeing “a minimum bandwidth” as recited in claim 9. *Guha* provides very little discussion of the functions of this component since the disclosure is directed to the QoS Enforcer that is at the head of the data center and not within the storage area network.

Choudhury fails to cure the identified deficiencies of *Guha*. *Choudhury* is directed generally to providing “different grades of service to customers sharing a resource” and providing “protection against overloads.” *Choudhury*, col. 6, ll. 1-3. *Choudhury* contains no teaching or suggestion of a “storage switch” that is “in communication with the at least one initiator and the at least one storage device” so that it can guarantee “a minimum bandwidth to the at least one initiator” and measure “the actual bandwidth of the at least one initiator,” as claim 9 recites (*emphasis added*). For example, *Choudhury* appears to use an “admission controller 203” to “deny a customer arrival access,” to

“determine the grade of access,” and to “guarantee minimum (GM) bounds.” *Id. at col. 9, l. 67-col. 10, l. 7.* There simply is nothing to teach or suggest a “storage switch” “in communication with at least one initiator and the at least one storage device” that performs the steps of claim 9 quoted above.

Accordingly, if *Choudhury* is combined with *Guha*, the resulting combination fails to teach or suggest these limitations of claim 9. Specifically, if the request and access control of *Choudhury* is combined with *Guha*, the functions would still be performed by a separate component at the head-end of a data center, or at best, by a “layer 4 switch” that is not in communication with a storage device. There is nothing within either reference, or their combination, that would suggest to one of ordinary skill in the art to utilize a “storage switch” in communication with both an initiator and storage device” to guarantee “a minimum bandwidth to the at least one initiator” and measure “the actual bandwidth of the at least one initiator,” as recited in claim 9.

Because *Guha* and *Choudhury*, either alone or in combination, fail to teach or suggest each of the limitations of claim 9, Applicants assert that this claim is patentable over the cited art. Claims 10-14 each ultimately depend from claim 9, and therefore, should be patentable for at least the same reasons.

Claims 16-21

Amended claim 16 recites similar limitations to those of claim 9 discussed above. For at least the same reasons as set forth with respect to claim 9, Applicants assert that claim 16 is patentable over *Guha* and *Choudhury*, either alone or in combination. Claims 17-24 each ultimately depend from claim 16 and therefore, should be patentable for at least the same reasons as claim 16.

Claims 22-24

Amended claim 22 recites:

the storage network including at least one initiator, at least one storage device, and a storage switch, wherein the at least one initiator and the at least one storage device are both in communication with the storage switch, the method comprising:

providing a connection from the at least one initiator to the at least one storage device via the storage switch in the storage network; and
adjusting, by the storage switch, the number of requests allowed the at

least one initiator to keep the bandwidth utilized by the at least one initiator within a specified range. *Emphasis added.*

As shown with respect to claims 1 and 9, the combination of *Guha* and *Choudhury* fails to teach or suggest a “storage switch” in communication with “at least one initiator and the at least one storage device” that performs such functions as providing “quality of service,” guaranteeing “a minimum bandwidth,” or measuring an “actual bandwidth.” Similarly, the combination fails to teach or suggest such a “storage switch” that adjusts, “the number of requests allowed the initiator” as recited in claim 22. Accordingly, for these and the reasons set forth with respect to claims 1 and 9, Applicants assert that claim 22 is patentable over the combination of *Guha* and *Choudhury*. Claims 23 and 24 each ultimately depend from claim 22 and therefore, should be patentable for at least the same reasons as claim 22.

Claims 38-44

Claim 38 recites similar limitations to those of claim 9. For at least the same reasons as set forth above, Applicants assert that claim 38 is patentable over the cited art. Claims 39-44 each ultimately depend from claim 38 and therefore, should be patentable for at least the same reasons.

Claims 3-5 and 6-8

Claims 3-5 and 6-8 each ultimately depend from claim 1. As set forth above, *Guha* fails to disclose each limitation of claim 1. Applicants further assert that *Guha* fails to teach or suggest each of the limitations of claim 1. *Choudhury* is not asserted to, nor does it teach or suggest each of the limitations of claim 1. The combination of *Guha* and *Choudhury* also fails to teach or suggest each of the limitations of claims 1. Accordingly, Applicants assert that claim 1 is patentable under 35 U.S.C. § 103(a) over *Guha* and *Choudhury*, either alone or in combination. Therefore, Applicants assert that claims 3-5 and 6-8 are patentable for at least the same reasons.

Claim 29

Claim 29 depends from claim 25. As set forth above, *Guha* fails to disclose each limitation of claim 25. Applicants further assert that *Guha* fails to teach or suggest each of the limitations of claim 25. *Choudhury* is not asserted to, nor does it teach or suggest each of the limitations of claim

25. The combination of *Guha* and *Choudhury* also fails to teach or suggest each of the limitations of claims 25. Accordingly, Applicants assert that claim 25 is patentable under 35 U.S.C. § 103(a) over *Guha* and *Choudhury*, either alone or in combination. Therefore, Applicants assert that claim 29 is patentable for at least the same reasons.

Claims 31-32

Claims 31-32 each ultimately depend from claim 30. As set forth above, *Guha* fails to disclose each limitation of claim 30. Applicants further assert that *Guha* fails to teach or suggest each of the limitations of claim 30. *Choudhury* is not asserted to, nor does it teach or suggest each of the limitations of claim 30. The combination of *Guha* and *Choudhury* also fails to teach or suggest each of the limitations of claims 30. Accordingly, Applicants assert that claim 30 is patentable under 35 U.S.C. § 103(a) over *Guha* and *Choudhury*, either alone or in combination. Therefore, Applicants assert that claims 31-32 are patentable for at least the same reasons.

Claim 34

Claim 34 depends from claim 33. As set forth above, *Guha* fails to disclose each limitation of claim 33. Applicants further assert that *Guha* fails to teach or suggest each of the limitations of claim 33. *Choudhury* is not asserted to, nor does it teach or suggest each of the limitations of claim 33. The combination of *Guha* and *Choudhury* also fails to teach or suggest each of the limitations of claims 33. Accordingly, Applicants assert that claim 33 is patentable under 35 U.S.C. § 103(a) over *Guha* and *Choudhury*, either alone or in combination. Therefore, Applicants assert that claim 34 is patentable for at least the same reasons.

V. Conclusion

Based on the above amendments and these remarks, reconsideration of Claims 1, 3-25, and 27-44 is respectfully requested.


The Examiner's prompt attention to this matter is greatly appreciated. Should further questions remain, the Examiner is invited to contact the undersigned attorney by telephone.

Enclosed is a PETITION FOR EXTENSION OF TIME UNDER 37 C.F.R. § 1.136 for extending the time to respond up to and including today, November 10, 2005.

The Commissioner is authorized to charge any underpayment or credit any overpayment to Deposit Account No. 501826 for any matter in connection with this response, including any fee for extension of time, which may be required.

Respectfully submitted,

Date: November 10, 2005

By: 
David E. Cromer
Reg. No. 54,768

VIERRA MAGEN MARCUS HARMON & DENIRO LLP
685 Market Street, Suite 540
San Francisco, CA 94105-4206
Telephone: (415) 369-9660
Facsimile: (415) 369-9665